

MEGCOM: Min-Energy Group COMMunication in Multi-hop Wireless Networks

Kai Han ^{*†}, Liu Xiang [†], Jun Luo[†], and Yang Liu [‡]

^{*}*School of Computer Science, Zhongyuan University of Technology, China*

[†]*School of Computer Engineering, Nanyang Technological University, Singapore*

[‡]*School of Information Sciences and Engineering, Henan University of Technology, China*
 {hankai, xi0001iu, junluo}@ntu.edu.sg, enjoyang@gmail.com

Abstract—Given the increasing demand from wireless applications, designing energy-efficient group communication protocols is of great importance to multi-hop wireless networks. A group communication session involves a set of member nodes, each of them needs to send a certain number of data packets to all other members. In this paper, we consider the problem of building a shared multicast tree spanning the member nodes such that the total energy consumption of a group communication session using the shared multicast tree is minimized. Since this problem was proven as NP-complete, we propose, under our Min-Energy Group COMMunication (MEGCOM) framework, three distributed approximation algorithms with provable approximation ratios. When the transmission power of each wireless node is fixed, our first two algorithms have the approximation ratios of $\mathcal{O}(\ln(\Delta + 1))$ and $\mathcal{O}(1)$, respectively, where Δ is the maximum node degree in the network. When the transmission power of each wireless node is adjustable, our third algorithm again delivers a constant approximation ratio. We also use extensive simulations to verify the practical performance of our algorithms.

I. INTRODUCTION

Group communication (or all-to-all multicasting) is a very important primitive for distributed systems, as a large body of applications including, among others, social networking, online meeting, network gaming, resource sharing, and data management are heavily relying on its service [1], [2]. In a traditional setting (e.g., the Internet), investigations on group communication focus on the system reliability in the face of network failures or group member changes [2]. As the recent booming of multi-hop wireless networks for various purposes (e.g., wireless mesh networks to allow Internet access in remote areas or wireless sensor networks for habitat monitoring) further demand group communication to work on top of such networks (e.g., [3]), group communication is facing a new challenge: energy constraints for wireless nodes. As it is generally believed that the continuous development of the wireless technology will result in more and more wireless applications making use of group communication primitive (e.g., mobile social networking and gaming), designing energy-efficient group communication protocols has become imperative.

One naive way for designing an energy-efficient group communication protocol is to employ the existing one-to-many multicast protocols. In other words, one may construct an energy-efficient one-to-many multicast tree for

each group member. While this approach may leverage on the abundant research results on building minimum-energy (one-to-many) multicast tree in wireless networks (e.g., [4]–[7]), it is not exactly practical, as it would lead to a huge computation and communication overhead in maintaining all these multicast trees. Therefore, a more practical way is to construct one tree spanning the group members as a shared multicast tree.

Surprisingly, although the minimum energy one-to-many multicasting problem has been studied extensively, there is little work on min-energy group communication in multi-hop wireless networks. Similar to the min-energy one-to-many multicasting problem, the min-energy shared tree multicasting problem is also NP-hard [8]. Moreover, the problem becomes more challenging because an optimal tree also depends on how many packets each group member is about to multicast. Liang *et al.* [8] attempted to tackle this problem by proposing several approximation algorithms. However, the resulting approximation ratios are far from satisfactory. For example, their approximation ratios are in the same order of $|M|$, where M is set of group members.

In this paper, we study the Min-Energy Group COMMunication (MEGCOM) framework in multi-hop wireless networks. In particular, we seek to minimize the total energy consumption of a group communication session relying on a shared multicast tree. As the induced Minimum-Energy All-to-All Multicasting (MEAAM) problem is NP-hard, we propose three distributed approximation algorithms with guaranteed approximation ratios for MEAAM, under both cases of fixed transmission power and of adjustable transmission power. Our algorithms significantly improve the best known results. Specifically:

- 1) When the transmission power of each wireless node is fixed, our first algorithm has an approximation ratio of $4\ln(\Delta + 1) + 7$, where Δ is the maximum node degree in the network.
- 2) When the transmission power of each wireless node is fixed, our second algorithm delivers a 13-approximation to MEAAM.
- 3) When the transmission power of each wireless node is adjustable, we prove the existence of a constant approximation algorithm for MEAAM, and we also show that a straightforward algorithm leads to a constant

approximation ratio of 145.

In the remaining of the paper, we first briefly review the literature in Sec. II. After formally defining the models and problem in Sec. III, we present and analyze our three algorithms in Sec. IV, V, and VI respectively. We also perform extensive simulations of our algorithm in Sec. VII, and we finally conclude our paper in Sec. VIII.

II. RELATED WORK

There is a large body of literature on group communication and multicasting, but, given the space limitation, we have to confine our discussions to those aiming at minimizing the energy consumption, but to leave out non-tree-based approaches such as [9], [10].

The minimum-energy one-to-many multicasting problem has been studied in [4]–[7]. Wieselthier *et al.* [4] consider a scenario where each node can adjust its transmission power continuously, and proposed several greedy heuristics for the minimum-power broadcast/multicast routing problems. Wan *et al.* [5] prove that the heuristics proposed by [4] have linear approximation ratios, and provide several approximation algorithms with constant approximation ratios for the minimum-energy multicasting problem based on the approximate minimum Steiner tree algorithm. Liang [6] consider a scenario in which each wireless node can adjust its transmission power in a discrete fashion, and the communication links are symmetric. They propose a centralized approximation algorithm with performance ratio $\mathcal{O}(\ln K)$ for building a minimum-energy multicasting tree, where K is the number of destination nodes in a one-to-many multicast session. Li *et al.* [7] consider a case where all nodes have a fixed transmission power and the communication links are asymmetric. They convert the minimum-energy multicasting problem to an instance of the Directed Steiner Tree problem, and presented several heuristics.

To the best of our knowledge, the only work that studies the minimum-energy group communication problem is [8]. In [8], Liang *et al.* propose to build a shared multicast tree such that the total energy consumption of realizing a group communication session using the shared tree is minimized. They prove that finding such a shared tree is a NP-complete problem, and used the approximate Steiner tree algorithm proposed by [11] to solve the problem. When the transmission power of each node is fixed, they prove that the approximate minimum Steiner tree has an approximation ratio of $2(|M| + 1)$, where M is the set of group members. When the transmission power of each node is adjustable, they prove that the approximate Steiner tree has an approximation ratio of $\Omega(8|M|)$. They also proposed a distributed approximation algorithm with an approximation ratio of $\Omega(4|M|^2)$.

III. PRELIMINARIES AND PROBLEM DEFINITION

A multi-hop wireless network is modeled by an undirected graph $G = (V, E)$, where V is the set of wireless nodes in the network and E is the set of wireless links. Each node in $v \in V$ has a unique identifier $v.id$. The nodes in V are all equipped with an omni-directional antenna. We assume that the transmission (tx) power of each node can be either identically fixed or continuously adjustable. When the tx power is fixed, we use ε_s to denote the energy consumption of transmitting a data packet by a node. When the tx power is adjustable, the energy required by any node u to transmit a data packet to another node v can be determined by the Euclidean distance between u and v . Following a very common formula, we define such energy consumption to be $d_{(u,v)}^\alpha$, where $d_{(u,v)}$ is the Euclidean distance between u and v and α is a constant (usually between 2 and 4). In either case, we assume that each link $(u, v) \in E$ is assigned a weight which is the amount of energy required for u to send a data packet to v (or vice versa). We also assume that the receiving (rx) power is always less than the tx power, and we denote by ε_r the energy consumption of receiving a data packet by a node.

In a group communication session, there exists a set of group members $M \subseteq V$, and each node $u \in M$ needs to send $p(u)$ data packets to all other nodes in $M \setminus \{u\}$. We denote by k the sum of the numbers of data packets originated from the group members, i.e., $k = \sum_{v \in M} p(v)$. As we explained in Sec. I, instead of building $|M|$ multicast trees originated from each node in M , building a shared multicast tree spanning the nodes in M to support the group communication session is more convenient in realistic settings. Therefore, in order to minimize the total energy consumption of the group communication session, we actually aim at solving the following Minimum-Energy All-to-All Multicasting (MEAAM) problem:

Find an *optimal shared multicast tree* T_{opt} such that the energy consumption of carrying all k packets over T_{opt} is minimized.

The hardness of this problem is immediate from [8]:

Proposition 1: The MEAAM problem is NP-complete.

For convenience of description, we define some other notations here. For any multicast tree T in G , we denote by $nd(T)$ the set of nodes in T , by $lv(T)$ the set of degree-one nodes in T , by $in(T)$ the set of internal nodes (nodes with degree more than one) in T , and by $\Psi(T)$ the total energy consumption of realizing a group communication session using T . For a node $u \in V$, we denote the set of neighboring nodes of u in G by $nb(u, G)$, and let $nb^+(u, G) = nb(u, G) \cup \{u\}$ and $nb2(u, G) = \bigcup_{v \in nb(u, G)} nb^+(v, G)$, where the latter is the two-hop neighbors of u . In the following, we present three distributed algorithms to construct approximate trees for T_{opt} , they respectively achieve a logarithmic approximation for fixed tx power, a constant

approximation for fixed tx power, and a constant approximation for adjustable tx power.

IV. MEGCOM-LFP: LOGARITHMIC APPROXIMATION FOR MEAAAM WITH FIXED TX POWER

In this section, we propose a distributed approximation algorithm for MEAAAM with fixed tx power. A brief introduction on the idea of the algorithm comes first, followed by detailed algorithm descriptions and a performance analysis. We also discuss how to maintain the shared group communication tree in the face of member joining and leaving.

A. Algorithm Principles

Unlike the common wisdom that directly uses a Steiner tree to approximate T_{opt} , our algorithm has two stages (with the first one having two sub-stages) to construct a T_A that approximates T_{opt} :

- S1-a: Identify a *buddy set* $B \subseteq V$ such that, for each $v \in B$, $nb(v, G) \cap M \neq \emptyset$.
- S1-b: Find a *guardian set* $C \subseteq B$ for M , such that $M \subseteq \bigcup_{v \in C} nb^+(v, G)$.
- S2 : Construct an approximate Steiner tree T_A to span the nodes in C .

As a result, T_A is the union of the member set M , the guardian set C (along with edges it uses), and the Steiner tree that spans C . We term each $v \in C$ a *guarding node* and a member $u \in M \wedge u \in nb(v, G)$ the *guarded member* of v . The tricky part is that we identify C by only searching among B that contains nodes having at least one neighbor in M , which excludes nodes in M that has no neighbor in M . As shown in Figure 1, all non-member nodes having at




Figure 1. Examples for the buddy set selection. The black nodes are group members in M , the grey nodes are the non-member nodes chosen for B . Here B contains every node having a concentric tx circle surrounding it; the solid lines indicate edges (or links) implied by those tx circles.

least one member in its neighborhood belong to B , whereas a member node belongs to B only if it has at least one other member in its neighborhood.

The motivation behind our algorithm is that a good approximation T_A should have more nodes in $lv(T_A)$ and less nodes in $in(T_A)$, as a node in $in(T_A)$ consume tx power for all packets and rx power for those not originated from it, whereas a node in $lv(T_A)$ consumes only tx power for packets originated from it and otherwise consume only

rx power. We use two examples in Figure 2 to illustrate the advantage of trees constructed by our algorithm and straightforward Steiner trees that span M . Therefore, we




Figure 2. Comparing our algorithm with Steiner tree based on two instances of MEAAAM. In the first example (a)–(c), the group members is $\{2, 3, 4, 5\}$. The numbers of data packets originated from nodes 2, 3, 4, and 5 are 100, 100, 1, and 1, respectively. The tx power of each node is 10. The reception power at each node is 1. The optimal Steiner tree in (b) leads to a total energy consumption of 6646, while our multicast tree (optimal for this case) in (c) gives 4848. In the second example (d)–(e), a Steiner-based heuristic (d) can result in 6 nodes in $in(T_A)$, whereas our multicast tree (e) only has 5 such nodes.

argue that using a Steiner-based heuristic does not give a proper approximation to T_{opt} . By far, the best-known centralized approximation algorithm for MEAAAM applies a Steiner-based heuristic [8], it only has an approximation ratio of $2(|M| + 1)$. The performance analysis of our algorithm in Sec. IV-C shows that our algorithm achieves a much better approximation ratio.

B. Algorithm Details

As explained in Sec. IV-A, our algorithm has two stages. In fact, each node can be in different states in each stage, so we let each node $v \in V$ to maintain a variable $v.state$ in order to keep track of the algorithm progress. Initially, every node starts with a state **Inactive**, and the algorithm terminates when all members in M are **Treed** (i.e., join the shared group communication tree) and all **Treed** non-members are spanned by shortest paths to form an approximate Steiner tree. We show the finite state machine for each node in Figure 3. The state **Competing** is involved in the guardian




Figure 3. The finite state machine for each node.

set identification (S1), while the two states **Connecting** and **Merging** are for approximate Steiner tree construction (S2).

We first present the pseudocodes of our algorithm by **Algorithm 1**, then we explain the algorithm execution in details. As we do not assume that the distributed tree

Algorithm 1: Finding an approximate multicast tree T_A for the MEAM problem

Input: The network $G = (V, E)$ and the member set M . For each node $v \in V$, $v.state = \text{Inactive}$

Output: An approximate multicast tree T_A : all the **Treed** nodes and all the marked paths/edges

```

1 task SCAN() /* executed periodically */
2 if  $v.state = \text{Inactive} \wedge \exists u \in nb(v, G) : u \in M$  then
3   forall the  $u \in nb(v, G) : u \in M$  do
4     if  $u.state = \text{Inactive}$  or Competing then
5        $v.S \leftarrow v.S \cup \{u\}$ ;  $v.ct \leftarrow v.ct + 1$ 
6   if  $v.ct > 0$  then  $v.state \leftarrow \text{Competing}$ ;
      SENDCOMPETE( $v.id, v.ct$ ) to  $\forall u \in nb_2(v, G)$ 
7 else if  $v.state = \text{Connecting}$  then
8   forall the  $u \in M$  do
9     if  $u.state = \text{Treed}$  or Connecting then
10       $v.S \leftarrow v.S \cup \{u.gd\}$ ;  $v.ct \leftarrow v.ct + 1$ 
11       $v.Path \leftarrow v.Path \cup SP(v \rightsquigarrow u.gd)$ 
12   if  $v.ct = |M|$  then
13      $v.prefer \leftarrow \arg \min_u \{SP(v \rightsquigarrow u) \in v.Path\}$ 
14     SENDCNTREQ( $v.id, v.ld$ ) to  $v.prefer$ 
15 else if  $v.state = \text{Treed} \wedge v.gd = v$  then DISCOVER()
16 else if  $v.state = \text{Merging}$  then
17   if  $\neg \text{DISCOVER}()$  then  $v.state \leftarrow \text{Treed}$ 
18   else SENDCNTREQ( $v.id, v.ld$ ) to  $v.prefer$ 
19 upon RECVCOMPETE( $u.id, u.ct$ ) from  $\forall u \in nb_2(v, G)$ 
20 if  $v.ct = \max_{u \in nb_2(v, G) \cup \{v\}} \{u.ct\}$  then
21    $v.state \leftarrow \text{Connecting}$ ;  $v.gd \leftarrow v$ ;  $v.ld \leftarrow v$ 
22   SENDDECISION( $v.id$ ) to  $\forall u \in v.S$ 
23 else  $v.state \leftarrow \text{Inactive}$ 
24  $v.S \leftarrow \emptyset$ ;  $v.ct \leftarrow 0$ 
25 upon RECVDECISION( $u.id$ )
26  $v.state \leftarrow \text{Treed}$ ;  $v.gd \leftarrow u$ ; mark the edge to  $u$ 
27 upon REVCNTREQ( $u.id, u.ld$ )
28 if  $v.prefer = u$  then
29   mark the path suggested by the request
30   SENDACCEPT( $v.id, \min\{v.ld, u.ld\}$ ) to  $u$ 
31   if  $v.ld < u.ld$  then  $v.state \leftarrow \text{Merging}$ 
32   else  $v.state \leftarrow \text{Treed}$ ;  $v.ld \leftarrow u.ld$ 
33 upon RECVACCEPT( $u.id, u.ld$ )
34 mark the accepted path
35 if  $v.ld < u.ld$  then  $v.state \leftarrow \text{Merging}$ 
36 else  $v.state \leftarrow \text{Treed}$ ;  $v.ld \leftarrow u.ld$ 

```

construction process can start synchronously at all nodes, we set a periodically executed *task* SCAN() (lines 1 to 18)

to drive the process. This task includes two part: lines 2 to 6 take care of identifying a guardian set C in a **greedy** manner, while other codes are for Steiner tree. In addition, there are four *procedures* that respond to different events.

At the beginning of the guardian set identification stage (S1), every node having at least one member neighbor implicitly joins the **buddy set** B by changing its state to **Competing** (line 6), as it has the right to compete for acting as a guarding node. As the guarded (member) sets do not intersect if two buddy nodes are more than two hops way from each other, it is sufficient to confine the competitions with the two-hop range for each node (line 6). The node that guards the most wins (line 20); it hence changes its state (thus enter the second stage, line 21) and also notifies the guarded members to change their states (lines 22 and 23). As a result, these guarded members will not be counted in later competition (line 4). A node v uses $v.gd$ to record its guarding node. For the guarded members, the algorithm has already terminated. At the end of S1, all nodes in **Connecting** state implicitly compose the **guardian set** C .

A guarding node that enters the Steiner tree construction stage (S2) will periodically check whether other members are guarded and, for an already guarded one, get the shortest paths to its guarding node (lines 8 to 11). If the targeted multi-hop wireless network has a proactive routing protocol (e.g., OLSR [12]), then a node only needs to check its local routing table. Otherwise routing queries may need to be sent if the network uses a reactive routing protocol (e.g., AODV [13] DSR or [14]). The tree construction starts after all members are either guarded or enters S2. Our algorithm is similar to the Kruskal-based shortest path heuristic proposed in [15]. The basic idea is to construct a minimum spanning tree over the extended graph $G'(V, E')$, where an edge $(u, v) \in E'$ has a length of $SP(u \rightsquigarrow v)$ (the shortest path in G between u and v). As this heuristic is a distributed implementation of the 2-approximation algorithm proposed in [11], the resulting approximation ratio is indeed 2.

The second stage (S2) starts with all the nodes in **Connecting** state, and it gradually finds a tree that spans these nodes. In order to eventually lead to a minimum spanning tree over $G'(V, E')$, only the shortest edge in E' should be added each time. This can be easily distributed initially (when a node is not connected to any other node yet), as the local information on the shortest paths from a node is sufficient for it to decide which edge in E' is to be added (line 13). However, after some nodes are connected to form several fragments, the local information is not sufficient anymore, hence we need a leader (the node in state **Merging**, chosen by the smallest *id* principle in lines 31 and 35, and recorded by $v.ld$) and certain consensus for obtaining a common preferred node to connect to. This consensus is implemented in the procedure DISCOVER(). Essentially, this procedure returns true if such a preferred node can be identified, then the leader (on behalf of the current fragment)

sends a connection request to that node (line 18). Otherwise the algorithm terminates, as the approximate Steiner tree has been constructed (line 17).¹ It is important to note that a connection request is accepted only if both ends consider each other as preferred (lines 14, 18, and 28); this is meant to avoid producing loops.

C. Performance Analysis

We examine the approximation ratio and the complexity of our algorithm in this section. In particular, we show that T_A resulting from **Algorithm 1** is a $[4\ln(\Delta + 1) + 7]$ -approximation to T_{opt} , where Δ is the maximum node degree of G . Our result is an exponential improvement to the best known result [8]. We also show that the complexity of **Algorithm 1** is polynomial in $|V|$.

1) *Approximation Ratio*: Our proof for the approximation ratio is presented in three steps. In **Lemma 1**, we provide an upper bound of the number of the internal nodes in T_A . In **Lemma 2**, we provide two lower bounds of the total energy consumption of any multicast tree for the MEAM problem. These bounds are used by **Theorem 1** to finally show the approximation ratio of **Algorithm 1**.

Lemma 1: Let T_I be the multicast tree spanning the nodes in M such that $|in(T_I)|$ is minimized, we have: $|in(T_A)| \leq [4\ln(\Delta + 1) + 6] \cdot |in(T_I)|$.

Proof: Let T_S^* and T_S , respectively, be the minimum Steiner tree and the approximate Steiner tree (produced by S2 of **Algorithm 1**) that span the nodes in C . Given the fact that S2 of **Algorithm 1** is a distributed implementation of [11], T_S is a 2-approximation of T_S^* , so $|nd(T_S)| \leq 2|nd(T_S^*)|$. As $|in(T_A)| = |nd(T_S)|$ given the way T_A is constructed, we have:

$$|in(T_A)| \leq 2|nd(T_S^*)|. \quad (1)$$

Let T_I^- be the tree constructed from T_I by pruning all the degree-one nodes in T_I (we assume $|nd(T_I^-)| \neq 0$, otherwise, the proof becomes trivial). For any node $u \in C \setminus nd(T_I^-)$, we can find a node $u_1 \in M$ such that u is adjacent to u_1 . If u_1 is not in T_I^- , then there must exist a node u_2 in T_I^- such that u_2 is adjacent to u_1 . In other words, any node in $C \setminus nd(T_I^-)$ can be connected to T_I^- by a path whose length is no more than 2. Therefore, we have:

$$|nd(T_S^*)| \leq |nd(T_I^-)| + 2|C| = |in(T_I)| + 2|C|. \quad (2)$$

For any node $v \in B$, we define $\sigma(v) = nb^+(v, G) \cap M$. Clearly, we have $|\sigma(v)| \leq \Delta + 1$. Let $C^* \subseteq B$ be the guardian set in G that contains the minimum number of nodes. Since S1 of **Algorithm 1** implements the greedy set covering algorithm proposed in [16], we have:

$$\begin{aligned} |C| &\leq \left\lceil \ln \left(\max_{v \in B} |\sigma(v)| \right) + 1 \right\rceil \cdot |C^*| \\ &\leq [\ln(\Delta + 1) + 1] \cdot |C^*|. \end{aligned} \quad (3)$$

¹Actually, there are some details on pruning overlapping edges of different shortest paths, but we omit them for brevity.

Note that $nd(T_I^-) \cap B$ is also a guardian set. so

$$|C^*| \leq |nd(T_I^-)| = |in(T_I)|. \quad (4)$$

Summarizing all the results obtained by far, we have:

$$\begin{aligned} |in(T_A)| &\leq 2(|in(T_I)| + 2|C|) \\ &\leq 2|in(T_I)| + 4[\ln(\Delta + 1) + 1] \cdot |in(T_I)| \\ &= [4\ln(\Delta + 1) + 6] \cdot |in(T_I)|, \end{aligned}$$

this bounds $|in(T_A)|$ from above by $|in(T_I)|$. ■

Lemma 2: For any multicast tree T in G spanning the nodes in M , we have:

- 1) $\Psi(T) \geq k \cdot [\varepsilon_s + (|M| - 1) \cdot \varepsilon_r]$
- 2) $\Psi(T) \geq |in(T)| \cdot (\varepsilon_s + \varepsilon_r) \cdot k$

Proof: Each node $u \in M$ must transmit its own $p(u)$ data packets. Therefore, the total energy consumption for transmitting data packets in a group communication session is at least $\sum_{u \in M} p(u) \cdot \varepsilon_s = k \cdot \varepsilon_s$. Moreover, the total energy consumption for receiving the data packets originated from u is $p(u) \cdot (|nd(T)| - 1) \cdot \varepsilon_r$. Therefore,

$$\begin{aligned} \Psi(T) &\geq k \cdot \varepsilon_s + \sum_{u \in M} p(u) \cdot (|nd(T)| - 1) \cdot \varepsilon_r \\ &\geq k \cdot [\varepsilon_s + (|M| - 1) \cdot \varepsilon_r]. \end{aligned}$$

The total energy consumption for realizing a group communication session using T can also be written as:

$$\begin{aligned} \Psi(T) &= \sum_{u \in lw(T)} p(u) \cdot (|in(T)| + 1) \cdot \varepsilon_s + \\ &\quad \sum_{u \in in(T) \cap M} p(u) \cdot |in(T)| \cdot \varepsilon_s + \\ &\quad k \cdot (|nd(T)| - 1) \cdot \varepsilon_r \\ &= k \cdot |in(T)| \cdot \varepsilon_s + \sum_{u \in lw(T)} p(u) \cdot \varepsilon_s + \\ &\quad k \cdot (|nd(T)| - 1) \cdot \varepsilon_r. \end{aligned} \quad (5)$$

Therefore:

$$\begin{aligned} \Psi(T) &\geq k \cdot |in(T)| \cdot \varepsilon_s + k \cdot (|in(T)| + |lw(T)| - 1) \cdot \varepsilon_r \\ &\geq |in(T)| \cdot (\varepsilon_s + \varepsilon_r) \cdot k. \end{aligned}$$

These give the two claimed lower bounds on $\Psi(T)$. ■

Theorem 1: The multicast tree T_A constructed by **Algorithm 1** has an approximation ratio of $4\ln(\Delta + 1) + 7$ for the MEAM problem.

Proof: As $|lw(T_A)| \leq |M|$, using (5) we can get:

$$\begin{aligned} \Psi(T_A) &= k \cdot |in(T_A)| \cdot \varepsilon_s + \sum_{u \in lw(T_A)} p(u) \cdot \varepsilon_s + \\ &\quad k \cdot (|nd(T_A)| - 1) \cdot \varepsilon_r \\ &\leq k \cdot |in(T_A)| \cdot \varepsilon_s + \sum_{u \in M} p(u) \cdot \varepsilon_s + \\ &\quad k \cdot (|in(T_A)| + |M| - 1) \cdot \varepsilon_r \\ &\leq k \cdot |in(T_A)| \cdot (\varepsilon_s + \varepsilon_r) + \\ &\quad k \cdot [\varepsilon_s + (|M| - 1) \cdot \varepsilon_r]. \end{aligned} \quad (6)$$

Let $\kappa = \Delta + 1$. Using **Lemmas 1, 2** and (6), we have:

$$\begin{aligned}\Psi(T_A) &\leq (4\ln\kappa + 6) \cdot |in(T_I)| \cdot (\varepsilon_s + \varepsilon_r) \cdot k + \Psi(T_{opt}) \\ &\leq (4\ln\kappa + 6) \cdot |in(T_{opt})| \cdot (\varepsilon_s + \varepsilon_r) \cdot k + \Psi(T_{opt}) \\ &\leq (4\ln\kappa + 6) \cdot \Psi(T_{opt}) + \Psi(T_{opt}) \\ &= [4\ln(\Delta + 1) + 7] \cdot \Psi(T_{opt}),\end{aligned}$$

hence the claimed approximation ratio. \blacksquare

2) **Algorithm Complexity:** The first stage of **Algorithm 1** has a time complexity of $\mathcal{O}(\Delta)$, as for every $v : nb(v, G) \cap M \neq \emptyset$, at least one node in $nb(v, G)$ leaves **Inactive** state every round.² In each round, all nodes in $nb2(v, G)$ sends a message, so the message complexity is $\mathcal{O}(\Delta \cdot |V|)$, where we use $|V|$ to bound from above the cardinality of $nb2(v, G)$. The second stage is a distributed Steiner tree algorithm. Its worst case time complexity and message complexity are $\mathcal{O}(D(G) \cdot |M|)$ and $\mathcal{O}(|M| \cdot |V|)$, respectively, as (i) only one group member leaves the **Connecting** state in each round in the worst case (hence $|M|$ rounds in total), (ii) differing from the first stage where the length of each round is a constant, this time the length of a round is determined by the round-trip time of a path (hence bounded by $D(G)$, the diameter of G), and (iii) the number of messages in each round is bounded by $|V|$.

D. Tree Maintenance

We discuss the tree maintenance with respect to the two individual stages separately. Actually, maintaining the shared tree simply for allowing the group communication session to continue is a separate issue that has been tackled in the literature (e.g., [17]–[19]). Therefore, our focus is only on how to maintain the proven approximation ratio of the tree.

When a member v joins or leaves the group communication session, if it has no existing guarding node in $nb(v, G)$ (for join) or is the only guarded node (for leave), then the impact directly goes to the Steiner tree (which will be discussed later). Otherwise v could directly join or leave in a localized manner by notifying its guarding node. The complication comes when a joining or leaving violates the greedy cover principle (thus affecting the approximation ratio). Therefore, we may need to re-execute S1 of **Algorithm 1** under those circumstances to maintain the greedy cover, hence maintain the proven approximation ratio.

If a member joining or leaving does not lead to any changes in the existing guardian set C , then the approximate Steiner tree remains intact. If a member joining brings one more node to C , we may need to re-execute S2 of **Algorithm 1** in the worst cases (when this newly joined guarding node results in several short cuts). Fortunately, we also have another choice of having this node directly connected to a closest (in terms of shortest path) guarding node. According to [20], this will lead to a $\mathcal{O}(\log |M|)$

approximation to Steiner tree, and in turn a $\mathcal{O}(\ln |M| \cdot \ln \Delta)$ approximation to MEAAM. If a member leaving removes a node from C , it may partition the tree into several fragments. The existing procedures of **Algorithm 1** can take care of this case as far as a new leader is elected for each fragment, as this is exactly the same as the fragment merging phase of the Steiner tree construction stage.

V. MEGCOM-CFP: CONSTANT APPROXIMATION FOR MEAAM WITH FIXED TX POWER

The algorithm proposed in Sec. IV has a relatively high message complexity as we involve non-members in both stages. In this section, we propose a simplified algorithm that only involve group member in the first stage. Interestingly, we can show that this simplified algorithm has a constant approximation ratio.

A. The Algorithm

As this new algorithm share the same second stage (S2) with **Algorithm 1**, we only show the updated first stage (S1) in **Algorithm 2** and then explain details accordingly.

Algorithm 2: Finding an approximate multicast tree T'_A for the MEAAM problem

Input: $G = (V, E)$ and M . For each node $v \in V$,
 $v.state = \text{Inactive}$; $v.nb = nb(v, G)$

Output: An approximate multicast tree T'_A

```

1 task SCAN() /* executed periodically */
2 if  $v.state = \text{Inactive} \wedge v \in M$  then
3   | SENDCOMPETE( $v.id$ ) to  $\forall u \in v.nb$ 
4 upon RECVCOMPETE( $u.id$ )
5 if  $v.state = \text{Inactive}$  then
6   | if  $v.id > u.id$  then SENDACCEPT( $v.id$ ) to  $u$ 
7 else SENDACK( $v.id, v.state$ )
8 upon RECVACCEPT( $u.id$ ) from  $\forall u \in v.nb$ 
9    $v.state \leftarrow \text{Connecting}$ ;  $v.gd \leftarrow v$ 
10 upon RECVACK( $u.id, u.state$ )
11 if  $u.state = \text{Connecting}$  then
12   |  $v.state \leftarrow \text{Treed}$ ;  $v.gd \leftarrow u$ ; mark the edge to  $u$ 
13 else if  $u.state = \text{Treed}$  then  $v.nb \leftarrow v.nb \setminus \{u\}$ 
```

The objective is again to identify a guardian set C' for M , but the algorithm differs from **Algorithm 1** in that C' is construct from M , i.e., $C' \subseteq M$. Similar to **Algorithm 1**, all the nodes that eventually change to **Connecting** state are in C' . The competition is again based on the smallest id principle. A node in **Inactive** state accepts any neighbor to be its guardian if its own id is larger (lines 5 to 6). If a node receives the acceptance from all its **Inactive** neighbors, it joins the guardian set C' (lines 8 to 9) by changing its state to **Connecting**. A node that is not in **Inactive** state always acknowledge the competition message by announcing its

²Our algorithm works in asynchronous settings; the concept of *round* is used only for evaluating the time complexity.

current state (line 7). A node either becomes guarded if it receives an acknowledgement that carries a **Connecting** state (line 12) or updates its inactive neighbor table if the state carried in the acknowledgement is **Treed**. Note that this algorithm only uses one-hop communications and finds an arbitrary guardian set (sufficient to achieve a constant approximation ratio), but we could also apply a similar procedure in **Algorithm 1** to find a greedy guardian set, which needs two-hop communications.

B. Performance Analysis

The approximation ratio of T'_A obtained by **Algorithm 2** is immediate from the following theorem.

Theorem 2: $\Psi(T'_A) \leq 13\Psi(T_{opt})$

Proof: Let T_I again be the multicast tree spanning the nodes in M such that $|in(T_I)|$ is minimized. The outcome of **Algorithm 2** implies that the nodes in C' are mutually independent. Therefore, any node in $in(T_I)$ can be adjacent to at most 5 nodes in C' [21]. Moreover, as any node $u \in C'$ is a member of M , u is either in $in(T_I)$ or is adjacent to certain node in $in(T_I)$. So we can get:

$$|C'| \leq 5|in(T_I)| \quad (7)$$

Let the minimum Steiner tree and approximate Steiner tree (produced by S2 of **Algorithm 2**) spanning the nodes in C' be \hat{T}_S^* and \hat{T}_S , respectively. We have:

$$|in(T'_A)| = |nd(\hat{T}_S)| \leq 2|nd(\hat{T}_S^*)| \quad (8)$$

Since any nodes in C' is either in $in(T_I)$ or is adjacent to certain node in $in(T_I)$, we can find a Steiner tree spanning the nodes in C' whose number of nodes is at most $|C'| + |in(T_I)|$. Therefore, we have:

$$|nd(\hat{T}_S^*)| \leq |C'| + |in(T_I)| \quad (9)$$

Combining (7)–(9), we can get:

$$|in(T'_A)| \leq 2(|C'| + |in(T_I)|) \leq 12|in(T_I)| \quad (10)$$

The result by far gives an upper bound for $|in(T'_A)|$. Now we can apply **Lemma 2** and (6) (similar to the proof of **Theorem 1**) to obtain $\Psi(T'_A) \leq 13\Psi(T_{opt})$. ■

We omit the complexity analysis for this algorithm, as it is very similar to **Algorithm 1**, except that the constant length of a round in the first stage becomes shorter, as only one-hop communications are required. The tree maintenance becomes simpler, as the guardian set is not constructed in a greedy manner, the impact of member joining or leaving directly goes to the Steiner tree.

VI. MEGCOM-CAP: CONSTANT APPROXIMATION FOR MEAM WITH ADJUSTABLE TX POWER

When the tx power of each node is adjustable, the idea of guardian set does not work anymore, as whether a node can be guarded by another is not known before an actual tx power is chosen. Even if we have a shared

multicast tree for the group communication session, any node u in a multicast tree T may need to use the tx power $\max\{d_{(u,v)}^\alpha | v \in nb(u, T)\}$ to transmit the data packets. Therefore, different nodes may use different tx power to forward the data packets. This makes the MEAM problem more complicated.

To design an approximation algorithm for the MEAM problem in the adjustable-tx-power case, we use a constructive proof to show the existence of a constant approximation algorithm, which naturally suggests one possible algorithm to construct the multicast tree. We still try to build a multicast tree whose sum of tx power of the internal nodes is the minimized, because the internal nodes still have a heavier forwarding load than the degree-one nodes for any multicas tree spanning M . Let $\lambda(u, T) = \max\{d_{(u,v)}^\alpha | v \in nb(u, T)\}$, and let T_W be a multicast tree such that $\Theta(T_W) = \sum_{u \in in(T_W)} \lambda(u, T_W)$ is minimized. Our idea is to find a multicast tree that approximates T_W and in turn approximates T_{opt} . The trick is to identify a quantitative relation between Θ and Ψ . This is done by **Theorem 3**, which in turn rely on the upper bounds of the total energy consumption for transmitting and receiving data, respectively, derived in **Lemma 3** and **Lemma 4**.

Lemma 3: For any multicast tree T spanning the nodes in M , the total energy consumption of transmitting all data packets using T in a group communication session is at most $2k \cdot \Theta(T)$.

Proof: Each internal node in $in(T)$ must transmit every data packet sent by each member node. So the total energy consumption for transmitting data packets by the internal nodes in T is:

$$\sum_{u \in in(T)} \lambda(u, T) \cdot \sum_{v \in M} p(v) = k \cdot \Theta(T). \quad (11)$$

Any degree-one node in T has a unique neighboring node in T . Therefore, there exists a function $f : lv(T) \rightarrow in(T)$ such that $f(u)$ is the unique internal node adjacent to $u : \forall u \in lv(T)$. Let $Z = \{f(u) | u \in lv(T)\}$. Clearly, for any degree one node $u \in T$, we have $\lambda(u, T) \leq \lambda(f(u), T)$. The total energy consumption for transmitting data packets by the degree-one nodes in T is:

$$\begin{aligned} & \sum_{u \in lv(T)} \lambda(u, T) \cdot p(u) \\ &= \sum_{v \in Z} \sum_{u \in \{w | w \in lv(T), f(w)=v\}} \lambda(u, T) \cdot p(u) \\ &\leq \sum_{v \in Z} \sum_{u \in lv(T)} \lambda(v, T) \cdot p(u) \\ &\leq k \cdot \sum_{v \in Z} \lambda(v, T) \\ &\leq k \cdot \Theta(T). \end{aligned} \quad (12)$$

Combining (11) and (12), the lemma follows. ■

Lemma 4: For any multicast tree T spanning the nodes in M , the total energy consumption of receiving all data packets using T in a group communication session is at most $\Psi(T_{opt}) + k \cdot \Theta(T)$.

Proof: Based on our assumption that $\forall v \in T, \varepsilon_r \leq \lambda(v, T)$, we can get:

$$\begin{aligned} |in(T)| \cdot k \cdot \varepsilon_r &= k \cdot \sum_{v \in in(T)} \varepsilon_r \\ &\leq k \cdot \sum_{v \in in(T)} \lambda(v, T) \\ &= k \cdot \Theta(T). \end{aligned}$$

The total energy consumption for receiving all the data packets by the nodes in T is $k \cdot (|nd(T)| - 1) \cdot \varepsilon_r$, and as $|lv(T)| \leq |M| \leq |nd(T_{opt})|$, we have:

$$\begin{aligned} &k \cdot (|nd(T)| - 1) \cdot \varepsilon_r \\ &= k \cdot [|lv(T)| + |in(T)| - 1] \cdot \varepsilon_r \\ &\leq k \cdot [|nd(T_{opt})| - 1] \cdot \varepsilon_r + k \cdot |in(T)| \cdot \varepsilon_r \\ &\leq \Psi(T_{opt}) + k \cdot \Theta(T), \end{aligned} \quad (13)$$

hence the claimed upper bound follows. \blacksquare

Theorem 3: Let T_W be a multicast tree spanning the nodes in M such that $\Theta(T_W)$ is minimized. For any multicast tree T , if $\Theta(T) \leq \rho \cdot \Theta(T_W)$, then $\Psi(T) \leq (3\rho + 1) \cdot \Psi(T_{opt})$

Proof: With **Lemma 3** and **Lemma 4**, we have:

$$\begin{aligned} \Psi(T) &\leq 3k \cdot \Theta(T) + \Psi(T_{opt}) \\ &\leq 3\rho k \cdot \Theta(T_W) + \Psi(T_{opt}) \\ &\leq 3\rho k \cdot \Theta(T_{opt}) + \Psi(T_{opt}) \\ &\leq (3\rho + 1)\Psi(T_{opt}), \end{aligned} \quad (14)$$

hence the theorem follows. \blacksquare

Remark: This theorem is itself very important. It shows that, if we could find a tree T that approximates T_W within a constant ratio, T is also a constant approximation to T_{opt} .

In this paper, we only suggest a straightforward algorithm that achieves a constant approximation ratio. In fact, we will show that, if we apply the distributed Steiner tree algorithm given in S2 of **Algorithm 1** to span M directly, the resulting tree \tilde{T}_S approximates T_W within a constant ratio.

Lemma 5: Let \tilde{T}_S be a 2-approximation Steiner tree in G spanning M , we have: $\Theta(\tilde{T}_S) \leq 48\Theta(T_W)$.

Proof: For any multicast tree T spanning the nodes in M , we denote by $\zeta(T)$ the sum of the weights (tx power) of the edges in T . For any node u in $in(T_W)$, we denote by $T_W^{(u)}$ the Euclidean minimum spanning tree of the nodes in $\{u\} \cup nb(u, T_W)$. According to [5], we have:

$$\zeta(T_W^{(u)}) \leq c \cdot \lambda(u, T_W), \quad 6 \leq c \leq 12. \quad (15)$$

Let G' be the graph constructed by superposing the $T_W^{(u)}$'s for all $u \in in(T_W)$. Let T_z be an arbitrary spanning tree of G' . We can get:

$$\begin{aligned} \zeta(T_z) &\leq \sum_{u \in in(T_W)} \zeta(T_W^{(u)}) \\ &\leq c \cdot \sum_{u \in in(T_W)} \lambda(u, T_W) \\ &= c \cdot \Theta(T_W). \end{aligned} \quad (16)$$

Let \tilde{T}_S^* be the minimum Steiner tree in G spanning the nodes in M . We have:

$$\zeta(\tilde{T}_S) \leq 2\zeta(\tilde{T}_S^*) \leq 2\zeta(T_z). \quad (17)$$

Note that the weight of any edge in \tilde{T}_S can be counted at most twice in $\Theta(\tilde{T}_S)$, so

$$\Theta(\tilde{T}_S) \leq 2\zeta(\tilde{T}_S). \quad (18)$$

The claimed approximation ratio follows by combining all these inequalities. \blacksquare

Finally, using **Theorem 3** and **Lemma 5** we can easily get:

Theorem 4: Our distributed Steiner tree algorithm on M approximates the MEAAM problem under the adjustable tx power case with a constant approximation ratio of 145.

VII. SIMULATIONS

In this section, we present the simulation results of our algorithms. We implement a simulator using C++ and use it to study how the performance of different MEGCOM algorithms is affected by various network parameters such as the network size, the node density and the percentage of group members. In the simulations, the network nodes are randomly deployed in a square area according to a certain node density, where the node density is defined as 1 if $|V|$ nodes are deployed upon a $\sqrt{|V|} \times \sqrt{|V|}$ square. The group members are selected from V following a Bernoulli distribution, and the number of packets originated from each group member is chosen from a uniform distribution $\mathcal{U}(1, 100)$. For each setting of the network parameters, we generate 100 network instances and show the mean values of the simulation results. Since our MEGCOM-CAP algorithm for the adjustable transmission power case is the approximate Steiner tree algorithm, which has already been shown in experiments to be superior to other algorithms [8], we only study the performance of MEGCOM-LFP and MEGCOM-CFP under the fixed transmission power case with $\varepsilon_s = 200$ and $\varepsilon_r = 20$.

For comparison, we build different multicast trees for the MEAAM problem in the simulations, using MEGCOM-LFP (**Algorithm 1**), MEGCOM-CFP (**Algorithm 2**), the Shortest Path Tree (SPT) algorithm, and the Steiner Tree algorithm. In the SPT algorithm, a root node is randomly selected from the group members, and the shared multicast tree is constructed by aggregating the shortest paths from the root node to other group members. In the implementation of the Steiner Tree algorithm, we adopt the approximation algorithm proposed by [11]; this is the Steiner tree algorithm also used in [8] as the solution to the MEAAM problem, and used in [5] as the solution to the min-energy one-to-many multicasting problem. We refrain from implementing other multicasting algorithms for the MEAAM problem such as the algorithms proposed by [17], [18], since they have been shown to be inferior to the Steiner Tree algorithm [8].




Figure 4. Comparing different algorithms for building multicast trees with the network density fixed to be 2.

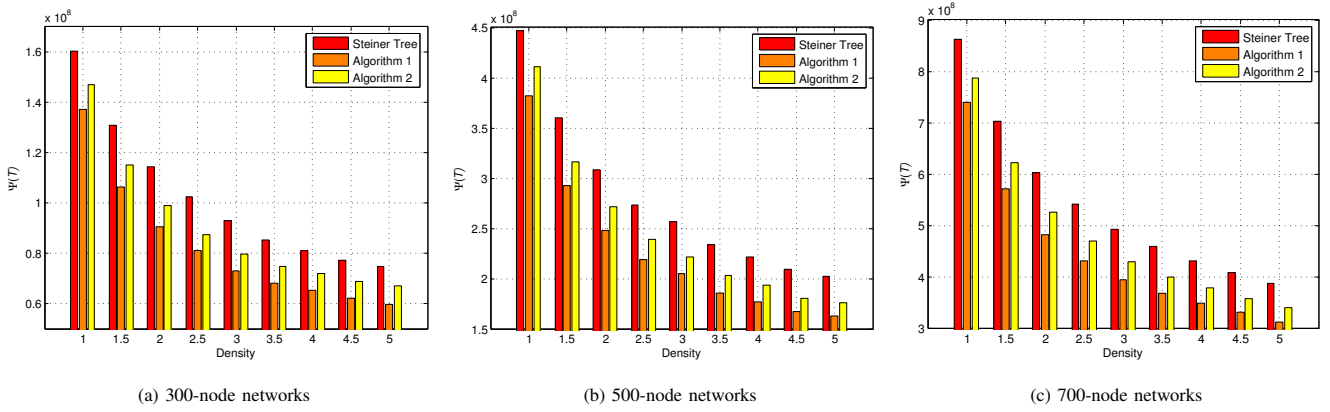


Figure 5. Comparing different algorithms for building multicast trees with the percentage of group members fixed to be 60%.

In Figure 4, the network density is fixed to 1 and the transmission range of each node is set to 2 to maintain the network connectivity. The network size is set to 300, 500, and 700 in Figure 4(a)-(c), respectively, and we vary the percentage of the group members from 10% to 90%. It can be seen from Figure 4 that the SPT algorithm always consumes the highest energy. This can be explained by the fact that the SPT tree is composed of shortest paths from one group member to the rest, introducing a large number of internal nodes that are not group members. Moreover, we can also see that **Algorithm 1** and **Algorithm 2** both outperform the Steiner tree algorithm. This phenomenon validates our analysis on approximation ratio in Section IV and Section V. The main reason for this phenomenon is that **Algorithm 1** and **Algorithm 2** both reduce the internal nodes of the shared multicast tree, whose energy consumption is the predominant part of the total energy consumption of a group communication session. Finally, we observe from Figure 4 that our algorithms perform better when the percentage of group members increases. Actually, when 90% network nodes become group members, **Algorithm 1** saves up to 40% energy cost compared with the SPT algorithm, or 25% compared with the Steiner tree algorithm. This can

be understood by the fact that more energy are conserved by our algorithms when more data packets are involved in the group communication session.

In Figure 5, we study the impact of node density on the performance of different algorithms. This time we remove the results for the SPT algorithm, as it perform worst in the previous comparisons. The percentage of group members is fixed to 60%, and the node density scales from 1 to 5 with an increment of 0.5. Again, the network size is set to 300, 500, and 700 in Figure 5(a)-(c), respectively. We observe that **Algorithm 1** and **Algorithm 2** always perform better than the Steiner tree algorithm in Figure 5. In particular, **Algorithm 1** saves 15%~20% energy compared with the Steiner tree algorithm. This demonstrates that the superiority of our MEGCOM algorithms persists when the node density changes. Here we use density as a rough indicator of Δ , as generating degree constrained networks is non-trivial. And even if we generated such graphs, they could hardly reflect any practical WSN deployments.

An interesting fact revealed by the simulations is that **Algorithm 1** always performs better than **Algorithm 2**, although the latter has a better theoretical approximation

ratio than the former (see *Theorem 1* and *Theorem 2*).³ We can attribute this effect to the fact that the guardian set is searched from a larger solution space in **Algorithm 1** than in **Algorithm 2**, which should on average result in less internal nodes in the shared multicast tree constructed by **Algorithm 1**, hence a lower total energy consumption of a group communication session using the tree and a better performance of **Algorithm 1** in average sense.

VIII. CONCLUSION

In this paper, we have studied the Minimum-Energy All-to-All Multicasting (MEAAM) problem in multi-hop wireless networks, where the transmission power of each wireless node could be either fixed or adjustable. Since the MEAAM problem is NP-complete, we have provided a set of distributed approximation algorithms under our MEGCOM (Min-Energy Group COMMunication) framework. For each algorithm, we have proven its approximation ratio with respect to the optimal solution. Our approximation ratios are significantly better than those of the best-known approximation algorithms for the MEAAM problem. We have further performed extensive simulations to validate our theoretical analysis and also to confirm the energy efficiency of our MEGCOM algorithms.

REFERENCES

- [1] D. Powell (Guest Ed.), "Group communication," *Comm. of the ACM*, vol. 39, pp. 50–97, 1996.
- [2] G. Chockler, I. Keidar, and R. Vitenberg, "Group communication specifications: A comprehensive study," *ACM Computing Surveys*, vol. 33, no. 4, p. 427469, 2001.
- [3] X. Yu and S. Chandra, "Designing an asynchronous group communication middleware for wireless users," in *Proc. of the 12th ACM MSWiM*, 2009, pp. 274–279.
- [4] J. Wieselthier, G. Nguyen, and A. Ephremides, "On the construction of energy-efficient broadcast and multicast trees in wireless networks," in *Proc. INFOCOM*, 2000, pp. 585–594.
- [5] P.-J. Wan, G. Calinescu, and C.-W. Yi, "Minimum-power multicast routing in static ad hoc wireless networks," *IEEE/ACM Trans. Netw.*, vol. 12, no. 3, pp. 507–514, 2004.
- [6] W. Liang, "Approximate minimum-energy multicasting in wireless ad hoc networks," *IEEE Trans. Mobile Comput.*, vol. 5, no. 4, pp. 377–387, 2006.
- [7] D. Li, Q. Liu, X. Hu, and X. Jia, "Energy efficient multicast routing in ad hoc wireless networks," *Computer Communications*, vol. 30, no. 18, pp. 3746–3756, 2007.
- [8] W. Liang, R. Brent, Y. Xu, and Q. Wang, "Minimum-energy all-to-all multicasting in wireless ad hoc networks," *IEEE Trans. Wireless Commun.*, vol. 8, no. 11, pp. 5490–5499, 2009.
- [9] S. J. Lee, W. Su, and M. Gerla, "On-demand multicast routing protocol in multihop wireless mobile networks," *Kluwer Mobile Networks & Applications*, vol. 7, no. 6, pp. 441–453, 2002.
- [10] B. Tavli and W. Heinzelman, "Energy-efficient real-time multicast routing in mobile ad hoc networks," *IEEE Trans. on Computers*, vol. 60, no. 5, pp. 707–722, 2011.
- [11] L. Kou, G. Markowsky, and L. Berman, "A fast algorithm for steiner trees," *Acta Informatica*, vol. 15, pp. 141–145, 1981.
- [12] T. Clausen and P. Jacquet, "Optimized Link State Routing Protocol (OSLR)," RFC 3626, October 2003, experimental.
- [13] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc On-Demand Distance Vector (AODV) Routing Protocol," RFC 3561, July 2003, experimental.
- [14] D. Johnson, Y. Hu, and D. Maltz, "The Dynamic Source Routing Protocol (DSR)," RFC 4728, February 2007, experimental.
- [15] F. Bauer and A. Varma, "Distributed algorithms for multicast path setup in data networks," *IEEE Trans. on Networking*, vol. 4, no. 2, pp. 181–191, 1996.
- [16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, Second Edition*. MIT Press and McGraw-Hill, 2001.
- [17] C.-C. Chiang, M. Gerla, and L. Zhang, "Adaptive shared tree multicast in mobile wireless networks," in *Proc. IEEE GLOBECOM*, vol. 3, 1998, pp. 1817–1822.
- [18] L. Ji and M. Corson, "A lightweight adaptive multicast algorithm," in *Proc. IEEE GLOBECOM*, vol. 2, 1998, pp. 1036–1042.
- [19] X. Jia, "A distributed algorithm of delay-bounded multicast routing for multimedia applications in wide area networks," *IEEE/ACM Trans. on Networking*, vol. 6, no. 6, pp. 828–837, 1998.
- [20] M. Imase and B. M. Waxman, "Dynamic steiner tree problem," *SIAM J. on Discrete Math.*, vol. 4, no. 3, pp. 369–384, 1991.
- [21] P.-J. Wan, K. M. Alzoubi, and O. Frieder, "Distributed construction of connected dominating set in wireless ad hoc networks," in *Proc. INFOCOM*, 2002, pp. 1597–1604.

³One generally considers a constant approximation ratio to be better than a logarithmic ratio. However, as an approximation ratio only serves as a metric to measure the worst case performance of an algorithm, the average performance of the algorithm in practice may not be well characterized by its approximation ratio.